

Min-Cost Multiflows in Node-Capacitated Undirected Networks

Maxim A. Babenko * Alexander V. Karzanov †

January 7, 2011

Abstract

We consider an undirected graph $G = (VG, EG)$ with a set $T \subseteq VG$ of terminals, and with nonnegative integer capacities $c(v)$ and costs $a(v)$ of nodes $v \in VG$. A path in G is a *T-path* if its ends are distinct terminals. By a *multiflow* we mean a function F assigning to each *T-path* P a nonnegative rational *weight* $F(P)$, and a multiflow is called *feasible* if the sum of weights of *T*-paths through each node v does not exceed $c(v)$. The *value* of F is the sum of weights $F(P)$, and the *cost* of F is the sum of $F(P)$ times the cost of P w.r.t. a , over all *T*-paths P .

Generalizing known results on edge-capacitated multiflows, we show that the problem of finding a minimum cost multiflow among the feasible multiflows of maximum possible value admits *half-integer* optimal primal and dual solutions. Moreover, we devise a strongly polynomial algorithm for finding such optimal solutions.

Keywords: minimum cost multiflow, bidirected graph, skew-symmetric graph

*Department of Mechanics and Mathematics, Moscow State University; Leninskie Gory, 119991 Moscow, Russia; email: max@adde.math.msu.su.

†Institute for System Analysis of the RAS; 9, Prospect 60 Let Oktyabrya, 117312 Moscow, Russia; email: sasha@cs.isa.ru

1 Introduction

1.1 Multiflows

For a function $\varphi: X \rightarrow \mathbb{R}_+$ and a subset $A \subseteq X$, we write $\varphi(A)$ to denote $\sum_{x \in A} \varphi(x)$. The *incidence vector* of A in \mathbb{R}^X is denoted by χ^A , i.e. $\chi^A(e)$ is 1 for $e \in A$ and 0 for $e \in X - A$ (usually X will be clear from the context). When A is a multiset, $\chi^A(e)$ denotes the number of occurrences of e in A .

In an undirected graph G , the sets of nodes and edges are denoted by VG and EG , respectively. When G is a directed graph, we speak of arcs rather than edges and write AG instead of EG . A similar notation is used for paths, cycles, and etc.

A *walk* in G is meant to be a sequence $(v_0, e_1, v_1, \dots, e_k, v_k)$, where each e_i is an edge (or arc) and v_{i-1}, v_i are its endnodes; when G is a digraph, e_i is directed from v_{i-1} to v_i . Edge-simple (or arc-simple) walks are called *paths*.

We consider an undirected graph G and a distinguished subset $T \subseteq VG$ of nodes, called *terminals*. Nodes in $VG - T$ are called *inner*. A *T-path* is a path P in G whose endnodes are distinct terminals; we usually assume that all the other nodes of P are inner. The set of T -paths is denoted by \mathcal{P} . A *multiflow* is a function $F: \mathcal{P} \rightarrow \mathbb{Q}_+$. Equivalently, one may think of F as a collection

$$(1.1) \quad \{(\alpha_1, P_1), \dots, (\alpha_n, P_m)\}$$

(for some m), where the P_i are T -paths and the α_i are nonnegative rationals, called *weights* of paths. Sometimes (e.g., in [IKN98]) such a multiflow F is called *free* to emphasize that *all* pairs of distinct terminals are allowed to be connected by flows. The *value* $\text{val}(F)$ of F is $\sum_P F(P)$. For a node v , define

$$(1.2) \quad \widehat{F}(v) := \sum (F(P): P \in \mathcal{P}, v \in VP);$$

the function \widehat{F} on VG is called the *(node) load function*. Let $c: VG \rightarrow \mathbb{Z}_+$ be a nonnegative integer function of *node capacities*. We say that F is *feasible* if $\widehat{F}(v) \leq c(v)$ for all $v \in VG$.

Suppose we are given, in addition, a function $a: VG \rightarrow \mathbb{Z}_+$ of *node costs*. Then the *cost* $a(F)$ of a multiflow F is the sum $\sum_P a(P)F(P)$, where $a(P)$ stands for the cost $a(VP)$ of a path P .

In this paper we consider the following problem:

(N) *Given G, T, c, a as above, find a multiflow F of minimum possible cost $a(F)$ among all feasible multiflows of maximum possible value.*

1.2 Previous results

When $|T| = 2$, (N) turns into the undirected min-cost max-flow problem under node capacities and costs, having a variety of applications; see, e.g., [FF62, La76]. It admits integer optimal primal and dual solutions [FF62].

In the special case $a \equiv 0$, we are looking simply for a feasible multiflow of maximum value. Such a problem has half-integer optimal primal and dual solutions, due to results

of Pap [Pa07] and Vazirani [Va01], respectively. Also it is shown in [Pa07] that the problem is solvable in strongly polynomial time by using the ellipsoid method.

An edge-capacitated version of (N) has been well studied. In this version, denoted by (E) , c and a are functions on EG rather than VG . For a multiflow F , its *edge load function* is defined similarly to (1.2):

$$(1.3) \quad \hat{F}(e) := \sum (F(P) : P \in \mathcal{P}, e \in EP) \quad \text{for all } e \in EG,$$

and its cost is defined to be $\sum_P a(EP)F(P)$. Problem (E) is reduced to (N) by adding an auxiliary node on each edge, but no converse reduction is known.

An old result is that (E) has a half-integer optimal solution [Ka79]. Also it is shown in [Ka94] that (E) has a half-integer optimal dual solution and that half-integer primal and dual optimal solutions can be found in strongly polynomial time by using the ellipsoid method. A “purely combinatorial” weakly polynomial algorithm, based on cost and capacity scaling, is devised in [GK97].

In the special case of (E) with $a \equiv 0$, the half-integrality results are due to Lovász [Lo76] and Cherkassky [Ch77], and a strongly polynomial combinatorial algorithm is given in [Ch77] (see also [IKN98] for faster algorithms).

1.3 New results

In this paper we prove that (N) always admits a half-integer optimal primal and dual solutions. In particular, this implies all half-integrality results mentioned in the previous subsection.

Similar to [Ka94], we introduce a *parametric* generalization of (N) , study properties of geodesics (shortest T -paths with respect to some length function), and reduce the parametric problem to a certain single-commodity flow problem. However, the details of this construction are more involved. In particular, the reduced problem concerns integer flows in a *bidirected* graph.

The second goal is to explore the complexity of (N) . We show that half-integer optimal primal and dual solutions to the parametric problem (and therefore to (N)) can be found in strongly polynomial time by using the ellipsoid method.

2 Preliminaries

2.1 Parametric problem and its dual

Instead of (N) , it is convenient to consider a more general problem, namely:

(N_λ) Given G, T, c, a as in (N) and, in addition, $\lambda \in \mathbb{Z}_+$, find a feasible multiflow F maximizing the objective function $\Phi(F, a, \lambda) := \lambda \cdot \text{val}(F) - a(F)$.

We will prove the following

Theorem 2.1 *For any $\lambda \in \mathbb{Z}_+$, problem (N_λ) has a half-integer optimal solution.*

(Note that $\Phi(F, qa, q\lambda) = q \cdot \Phi(F, a, \lambda)$ for any multiflow F and $q \in \mathbb{Q}_+$. Therefore, the optimality of a multiflow in the parametric problem preserves when both a and λ are

multiplied by the same positive factor q . This implies that the theorem is generalized to arbitrary $a : VG \rightarrow \mathbb{Q}_+$ and $\lambda \in \mathbb{Q}_+$ (but keeping the integrality of c). However, we prefer to deal with integer-valued a and λ in what follows.)

By standard linear programming arguments, (N) and (N_λ) become equivalent when λ is large enough (moreover, the existence of a half-integer optimal solution for (N_λ) easily implies that taking $\lambda := 2c(VG)a(VG) + 1$ is sufficient).

Problem (N_λ) can be viewed as a linear program with variables $F(P) \in \mathbb{Q}_+$ assigned to T -paths P . Assign to a node $v \in VG$ a variable $l(v) \in \mathbb{Q}_+$. Then the linear program dual to (N_λ) is:

(D_λ) *Minimize $c \cdot l$ provided that the following holds for every T -path P :*

$$(2.1) \quad l(P) \geq \lambda - a(P).$$

2.2 Translating to edge lengths

The above dual problem (D_λ) involves lengths of paths (namely, $l(P)$ and $a(P)$) determined by “lengths” of nodes (l and a , respectively). It is useful to transform lengths of nodes into lengths of edges. To do so, for $w : VG \rightarrow \mathbb{Q}_+$, we define the function \bar{w} on EG by

$$(2.2) \quad \bar{w}(e) := \alpha_u w(u) + \alpha_v w(v) \quad \text{for } e = uv \in EG,$$

where $\alpha_x := \frac{1}{2}$ if $x \in VG - T$, and $\alpha_x := 1$ if $x \in T$. This provides the correspondence

$$(2.3) \quad w(P) = \bar{w}(P) \quad \text{for each } T\text{-path } P$$

(where $w(P)$ stands for $w(VP)$, and $\bar{w}(P)$ for $\bar{w}(EP)$). For a, l as above, define

$$(2.4) \quad \ell := \bar{l} + \bar{a}.$$

Let $\text{dist}_\ell(u, v)$ denote the ℓ -distance between vertices u and v , i.e. the minimum ℓ -length $\ell(P)$ of a u - v path P in G . Then, in view of (2.3) and (2.4), the constraints in (D_λ) can be rewritten as

$$(2.5) \quad \text{dist}_\ell(s, t) \geq \lambda \quad \text{for all } s, t \in T, s \neq t.$$

By the linear programming duality theorem applied to (N_λ) and (D_λ) , a feasible multiflow F and a function $l : VG \rightarrow \mathbb{Q}_+$ satisfying (2.5) are optimal solutions to (N_λ) and (D_λ) , respectively, if and only if the following (complementary slackness conditions) hold:

(2.6) if P is a T -path and $F(P) > 0$, then $\ell(P) = \lambda$; in particular, P is ℓ -shortest;

(2.7) if $v \in VG$ and $l(v) > 0$, then v is *saturated* by F , i.e. $\widehat{F}(v) = c(v)$.

In the rest of the paper, to simplify technical details, we will always assume that the input costs a of all nodes are *strictly positive*. Then the edge lengths ℓ defined by (2.4) are strictly positive as well. This assumption will not lead to loss of generality in essence, since the desired results for a nonnegative input cost function a can be obtained by applying a perturbation technique in spirit of [Ka94, pp. 320–321] (by replacing a by an appropriate strictly positive cost function).

2.3 Geodesics

Condition (2.6) motivates the study of the structure of ℓ -shortest T -paths in G . To this aim, set $p := \min\{\text{dist}_\ell(s, t) \mid s, t \in T, s \neq t\}$. A T -path P such that $\ell(P) = p$ is called an ℓ -geodesic (or just geodesic if ℓ is clear from the context). When a multiflow F in G is given as a collection (1.1) in which all paths P_i are ℓ -geodesics, we say that F is an ℓ -geodesic multiflow.

Next we utilize one construction from [Ka79, Ka94], with minor changes. Consider a node $v \in VG$. Define the potential $\pi(v)$ to be the ℓ -distance from v to the nearest terminal, i.e. $\pi(v) := \min\{\text{dist}_\ell(v, t) \mid t \in T\}$. Set $VG_\ell := \{v \in VG \mid \pi(v) \leq \frac{1}{2}p\}$ (in particular, $T \subseteq VG_\ell$). For $s \in T$, define $V^s := \{v \in VG \mid \text{dist}_\ell(s, v) < \frac{1}{2}p\}$. Also define $V^\sharp := \{v \in VG \mid \pi(v) = \frac{1}{2}p\}$. We refer to V^s as the zone of a terminal $s \in T$, and to V^\sharp as the set of central nodes (w.r.t. ℓ). The sets V^s ($s \in T$) and V^\sharp are pairwise disjoint and give a partition of VG_ℓ .

The following subset of edges is of importance:

$$\begin{aligned} EG_\ell := & \{uv \in EG \mid \exists s \in T : u \in V^s, v \in V^s \cup V^\sharp, |\pi(u) - \pi(v)| = \ell(uv)\} \\ & \cup \{uv \in EG \mid \exists s, t \in T, s \neq t : u \in V^s, v \in V^t, \pi(u) + \pi(v) + \ell(uv) = p\}. \end{aligned}$$

One can see that the subgraph $G_\ell := (VG_\ell, EG_\ell)$ of G contains all ℓ -geodesics. Moreover, a straightforward examination shows that the structure of ℓ -geodesics possesses the properties as expressed in the following lemma (which is, in fact, a summary of Claims 1–3 from [Ka94] and uses the strict positivity of ℓ).

Lemma 2.2 *Let P be an ℓ -geodesic running from $s \in T$ to $t \in T$. Then P is contained in G_ℓ and exactly one of the following takes place:*

1. *P contains no central nodes and can be represented as the concatenation $P_1 \circ (u, e, v) \circ P_2$, where $u \in V^s$, $v \in V^t$, $s \neq t$, and $e \in EG_\ell$.*
2. *P contains exactly one central node $w \in V^\sharp$ and can be represented as the concatenation $P = P_1 \circ (u, e_1, w, e_2, v) \circ P_2$, where $u \in V^s$, $v \in V^t$, $s \neq t$, and $e_1, e_2 \in EG_\ell$.*

In both cases, parts P_1 and P_2 are contained in the induced subgraphs $G_\ell[V^s]$ and $G_\ell[V^t]$, respectively. The potentials π are strictly increasing as we traverse P_1 from s to u , and strictly decreasing as we traverse P_2 from v to t .

Conversely, any T -path in G_ℓ obeying the above properties is an ℓ -geodesic.

3 Primal half-integrality

3.1 Auxiliary bidirected graph

In this subsection we introduce an auxiliary bidirected graph, which will be the cornerstone of our approach both for proving half-integrality results and for providing a polynomial-time algorithm.

Given G , T , c , a and λ as above, let l be an optimal solution to (D_λ) . Form the edge lengths $\ell := \bar{l} + \bar{a}$, the potential π , the subgraph G_ℓ , and the sets V^s ($s \in T$) and V^\sharp , as in Subsection 2.3. One may assume that $p := \min\{\text{dist}_\ell(s, t) \mid s, t \in T, s \neq t\} = \lambda$ (since $p \geq \lambda$, by (2.5), and if $p > \lambda$ then $F = 0$ is an optimal solution to (N_λ) , by (2.6)).

For further needs, we reset $c := 2c$, making all node capacities even integers. Now our goal is to prove the existence of an *integer* optimal multiflow F in problem (N_λ) (which is equivalent to proving the half-integrality w.r.t. the initial c).

Recall that in a *bidirected* graph (or a *BD-graph* for short) edges of three types are allowed: a usual directed edge, or an *arc*, that leaves one node and enters another one; an edge directed *from both* of its ends; and an edge directed *to both* of its ends (cf. [EJ70, Sc03]). When both ends of an edge coincide, the edge becomes a loop. For our purposes we admit no loop entering and leaving its end node simultaneously. Sometimes, to specify the direction of an edge $e = uv$ at one or both of its ends, we will draw arrows above the corresponding node characters. For example, we may write \overrightarrow{uv} if e is directed from u to v (a usual arc), $\overleftarrow{u}\overleftarrow{v}$ if e leaves both u, v , $\overleftarrow{u}\overrightarrow{v}$ if e enters both u, v , and $\overrightarrow{u}v$ if e leaves u (and either leaves or enters v).

A *walk* in a BD-graph is an alternating sequence

$$P = (s = v_0, e_1, v_1, \dots, e_k, v_k = t)$$

of nodes and edges such that each edge e_i connects nodes v_{i-1} and v_i , and for $i = 1, \dots, k-1$, the edges e_i, e_{i+1} form a *transit pair* at v_i , which means that one of e_i, e_{i+1} enters and the other leaves v_i . As before, an edge-simple walk is referred to as a *path*.

Now we associate to G_ℓ a BD-graph H with edge capacities $c: EH \rightarrow \mathbb{Z}_+$, as follows (see Fig. 1 for an illustration). Each noncentral node $v \in VG_\ell - V^\natural$ generates two nodes v^1, v^2 in H . They are connected by edge (arc) e_v going from v^1 to v^2 and having the capacity equal to $c(v)$. We say that e_v *inherits the capacity* of the node v . For $s \in T$, the set $\overline{V}^s := \{v^1, v^2 \mid v \in V^s\}$ in H is called the *zone* of s , similar to V^s in G .

Consider an edge $e = uv \in EG_\ell$. Let $u, v \in V^s$ for some $s \in T$ and assume for definiteness that $\pi(u) < \pi(v)$ (note that $\ell(uv) > 0$ implies $\pi(u) \neq \pi(v)$; this is where the strict positivity of the cost function a is important). Then e generates in H an edge (arc) going from u^2 to v^1 , and we assign infinite capacity to it. (By “infinite capacity” we mean a sufficiently large positive integer.) Now let $u \in V^s$ and $v \in V^t$ for distinct $s, t \in T$. Then e generates an infinite capacity edge $\overrightarrow{u}^2\overleftarrow{v}^2$ (leaving both u^2 and v^2).

The transformation of central nodes is less straightforward. Each $w \in V^\natural$ generates in H a so-called *gadget*, denoted by Γ_w . It consists of $|T| + 1$ nodes; they correspond to w and the elements of T and are denoted as θ_w and $\theta_{w,s}$, $s \in T$. The edges of Γ_w are: a loop e_w leaving θ_w (twice) and, for each $s \in T$, an edge $e_{w,s}$ going from $\theta_{w,s}$ to θ_w , called the *s-leg* in the gadget. Each edge in Γ_w is endowed with the capacity equal to $c(w)$.

Each gadget Γ_w is connected to the remaining part of H as follows. For each edge of the form vw in G_ℓ , we know that $v \in V^s$ for some $s \in T$ (by the construction of G_ℓ). Then vw generates an infinite capacity edge (arc) going from v^2 to $\theta_{w,s}$.

Finally, we add to H an extra node q , regarding it as the *source*, and for each $s \in T$, draw an infinite capacity edge (arc) from q to s^1 .

The obtained BD-graph H captures information about the ℓ -geodesics in G . Namely, each ℓ -geodesic P going from s to t induces a unique closed q - q walk \overline{P} in H . The first and the last edges of \overline{P} are \overrightarrow{qs}^1 and $\overleftarrow{t}^1\overleftarrow{q}$, respectively. For a noncentral node v in P , \overline{P} traverses the edge $\overrightarrow{v}^1\overrightarrow{v}^2$. An edge $uv \in EP$ with $\pi(u) < \pi(v)$ inside a zone induces the edge $\overrightarrow{u}^2\overrightarrow{v}^1$ in \overline{P} . An edge $uv \in EP$ connecting different zones (if any) induces the edge $\overrightarrow{u}^2\overleftarrow{v}^2$ in \overline{P} . Finally, suppose P traverses a central node $w \in V^\natural$ and let $uw, wv \in EG_\ell$

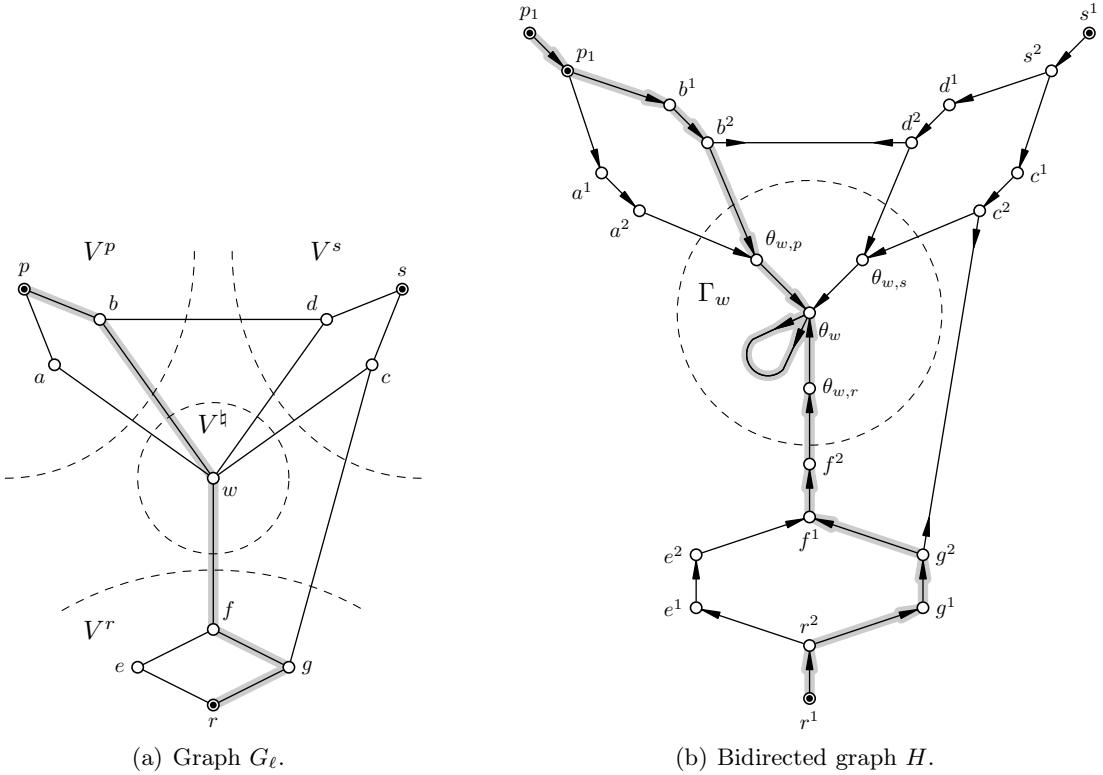


Figure 1: Constructing graph H . Here $T = \{p, s, r\}$, $V^p = \{p, a, b\}$, $V^s = \{s, c, d\}$, $V^r = \{r, e, f, g\}$, $V^h = \{w\}$. (The source q is not shown.) Bidirected edges leaving one endpoint and entering the other are indicated by ordinary directed arcs. Marked are one ℓ -geodesic P and its image \overline{P} .

be the edges of P incident to w . By Lemma 2.2, $u \in V^s$ and $v \in V^t$ for some $s \neq t$. Then the sequence of nodes u, w, v in P generates the subpath in \overline{P} with the sequence of edges $u^2 \theta_{w,s}, e_{w,s}, e_w, e_{w,t}, \theta_{w,t} v^2$.

The resulting walk \overline{P} is edge-simple, so it is a closed path. Conversely, let Q be a (nontrivial) q - q walk in H . One can see that Q with q removed is concatenated as $Q_1 \circ Q' \circ Q_2$, where Q_1 is a directed path within a zone \overline{V}^s , Q_2 is reverse to a directed path within a zone \overline{V}^t (with possibly $s = t$), and Q' either (i) is formed by an edge $\overrightarrow{u^2} \overleftarrow{v^2}$ connecting these zones (in which case $s \neq t$), or (ii) is the walk with the sequence of edges $e_{w,s}, e_w, e_{w,t}$, for some central node w of G_ℓ . Moreover, the image in G of each of Q_1, Q_2 is an ℓ -shortest path. When $s = t$ happens in case (ii), Q traverses the edge $e_{w,s}$ twice. In all other cases, Q is edge-simple and its image in G is an ℓ -shortest T -path (a λ -geodesic).

These observations show that there is a natural bijection between the ℓ -geodesics in G and the (nontrivial) q - q paths in H .

We will refer to the BD-graph H described above as the *compact* BD-graph related to G_ℓ ; it will be essentially used to devise an efficient algorithm for solving (N_λ) in Section 4. Besides, in the proof of the primal integrality (with c even) in Section 5, we will deal with

a modified BD-graph. It is obtained from H as above by replicating each gadget Γ_w into $c(w)$ copies Γ_{w^i} , $i = 1, \dots, c(w)$, called the *1-gadgets* generated by w . More precisely, to construct Γ_{w^i} , we make i -th copy θ_{w^i} of the node θ_w , i -th copy e_{w^i} of the loop e_w leaving θ_{w^i} (twice), and i -th copy $e_{w^i,s}$ of each leg $e_{w,s}$, $s \in T$, where $e_{w^i,s}$ goes from $\theta_{w,s}$ to θ_{w^i} (so $\theta_{w,s}$, $s \in T$, are the common nodes of the created 1-gadgets). All edges in these 1-gadgets are endowed with *unit* capacities.

We keep notation H for the constructed graph and call it the *expensive* BD-graph related to G_ℓ . Also we keep notation c for the edge capacities in H . There is a natural relationship between the q - q walks (paths) in both versions of H . The 1-gadgets created from the same central node w of G_ℓ are isomorphic, and for any $i, j = 1, \dots, c(w)$, there is an automorphism of H which swaps θ_{w^i} and θ_{w^j} and is invariant on the other nodes.

3.2 Bidirected flows

Let Γ be a bidirected graph. Like in usual digraphs, $\delta^{\text{in}}(v)$ and $\delta^{\text{out}}(v)$ denote the sets of edges in Γ entering and leaving $v \in V\Gamma$, respectively. A loop e at v , if any, is counted twice in $\delta^{\text{in}}(v)$ if e enters v , and twice in $\delta^{\text{out}}(v)$ if e leaves v ; hence $\delta^{\text{in}}(v)$ and $\delta^{\text{out}}(v)$ are actually multisets. (Recall that we do not allow a loop which simultaneously enters and leaves a node.)

Let q be a distinguished node with $\delta^{\text{in}}(q) = \emptyset$ in Γ (the *source*) and let the edges of Γ have integer capacities $c: E\Gamma \rightarrow \mathbb{Z}_+$. A *bidirected q -flow*, or a *BD-flow* for short, is a function $f: E\Gamma \rightarrow \mathbb{Q}_+$ satisfying $\text{div}_f(v) = 0$ for all nodes $v \in V\Gamma - \{q\}$; and the *value* of f is defined to be $\text{div}_f(q)$ (cf. [GK04]). Here

$$(3.1) \quad \text{div}_f(v) := f(\delta^{\text{out}}(v)) - f(\delta^{\text{in}}(v))$$

is the *divergence* of f at v . Note that if e is a loop at v then e contributes $\pm 2f(e)$ in $\text{div}_f(v)$. If $f(e) \leq c(e)$ for all $e \in E\Gamma$ then f is called *feasible*. In addition, if f is integer-valued on all edges then we refer to f as an integer bidirected q -flow, or an *IBD-flow*. One can see that finding a fractional (resp. integer) BD-flow of the maximum value is equivalent to constructing a maximum fractional (resp. integer) packing of closed q - q walks (they leave q twice).

Return to an optimal solution l to (D_λ) , and let $\ell := \bar{a} + \bar{l}$. Consider the (expensive or compact) BD-graph H related to G_ℓ , and the capacity function c on the edges of H (constructed from the node capacities c of G). The above correspondence between ℓ -geodesics in G and q - q paths in H is extended to ℓ -geodesic multiflows in G and certain q -flows in H (where q is the source in H as before). More precisely, let F be a (fractional) ℓ -geodesic multiflow in G represented by a collection of ℓ -geodesics P_i and weights $\alpha_i := F(P_i)$, $i = 1, \dots, m$ (cf. (1.1)). Then each P_i determines a q - q path \bar{P}_i in H , and $f := \alpha_1 \chi^{E\bar{P}_1} + \dots + \alpha_m \chi^{E\bar{P}_m}$ is a BD-flow in H ; we say that f is *generated* by F (note that $\text{val}(f) = 2 \text{val}(F)$). Furthermore, f is feasible if F is such, and for each central node $w \in V^\natural$, the following relations hold:

$$(3.2) \quad \sum_{s \in T} f(e_{w,s}) = 2f(e_w) \quad \text{and} \quad f(e_{w,s}) \leq f(e_w) \quad \text{for each } s \in T.$$

Considering an arbitrary BD-flow f in H , we say that f is *good* if it satisfies (3.2) for all $w \in V^\natural$ (here the second relation in (3.2) is important, while the first one obviously holds for any BD-flow). The following assertion is of use.

Lemma 3.1 *Let f be a good BD-flow in H . Then f is generated by an ℓ -geodesic multiflow F in G . Moreover, if f is integral, then it is generated by an integer ℓ -geodesic multiflow F . In both cases, F can be found in $O(|EH|)$ time.*

Proof. Suppose there is a central node $w \in V^\ddagger$ such that $f(e_w) > 0$. Let us say that $p \in T$ dominates at w (w.r.t. f) if $f(e_{w,p}) = f(e_w)$. From (3.2) it follows that there exist distinct $s, t \in T$ such that $f(e_{w,s}), f(e_{w,t}) > 0$ and none of $p \in T - \{s, t\}$ dominates at w . Choose such s, t . Build in H a maximal walk Q starting with $\theta_w, e_{w,s}, \theta_{w,s}, \dots$ and such that $f(e) > 0$ for all edges e of Q . It is easily seen from the construction of H that Q is edge-simple, terminates at q , and have all vertices in \overline{V}^s , except for $\theta_w, \theta_{w,s}$. Build a similar walk (path) Q' starting with $\theta_w, e_{w,t}, \theta_{w,t}$. Then the concatenation of the reverse to Q , the loop e_w and the path Q' is a q - q path and its image P in G is an ℓ -geodesic (from s to t).

Assign the weight of P to be the maximum number α subject to two conditions: (i) $\alpha \leq f(e)$ for each $e \in E\overline{P}$, and (ii) the flow $f' := f - \alpha\chi_{\overline{P}}$ is still good. If α is determined by (i), we have $|\text{supp}(f')| < |\text{supp}(f)|$ (where $\text{supp}(\varphi) := \{x \mid \varphi(x) \neq 0\}$), whereas if α is determined by (ii), there appears $p \in T$ dominating at w (w.r.t. f'). If $f'(e_w) > 0$, repeat the procedure for f' and w , otherwise apply the procedure to f' and another $w' \in V^\ddagger$, and so on. (Note that if, in the process of handling w , a current weight α is determined by (ii), then the weights of all subsequent paths through e_w are already determined by (i); this will provide the desired complexity.)

Eventually, we come to a good flow \tilde{f} with $\tilde{f}(e_w) = 0$ for all $w \in V^\ddagger$. This \tilde{f} is decomposed into a sum of flows along q - q paths in a straightforward way, in $O(|EH|)$ time (like for usual flows in digraphs). Taking together the images in G of the constructed weighted q - q paths, we obtain a required ℓ -geodesic multiflow F . The running time of the whole process is $O(|EH|)$, and if f is integral, then the weights α of all paths are integral as well. (Integrality of a current weight α subject to integrality of a current flow f is obvious when α is determined by (i), and follows from the fact, implied by (3.2), that for any $p \in T$, $\sum_{s \neq p} f(e_{w,s}) - f(e_{w,p})$ is even, when α is determined by (ii).) \square

Remark 3.2 *In case of an expensive BD-graph H , any feasible IBD-flow f is good. Indeed, for any 1-gadget Γ_{w^i} in H , we have $c(e_{w^i}) = 1$ and, therefore, $f(e_w) \in \{0, 1\}$. The second relation in (3.2) is trivial when $f(e_w) = 0$, and follows from the constraints $f(e_{w,s}) \leq c(e_{w,s}) = 1$ ($s \in T$) when $f(e_w) = 1$.*

Define the following subset of edges in H :

$$(3.3) \quad E_0 := \{e_v \mid v \in VG, \ l(v) > 0\}.$$

For an optimal (possibly fractional) solution F to (N_λ) and a node $v \in VG$ with $l(v) > 0$, we have $\hat{F}(v) = c(v)$ (by (2.7)); so the edge e_v of H corresponding to v is saturated by the BD-flow f generated by F , i.e. $f(e_v) = c(e_v)$. We call the edges in E_0 *locked*.

Thus, the graph H admits a (fractional) good feasible BD-flow saturating the locked edges. The following strengthening is crucial.

Proposition 3.3 *There exists a good feasible IBD-flow in H that saturates all locked edges.*

A proof of this proposition involves an additional graph-theoretic machinery and will be given in Section 5. Assuming its validity, we immediately obtain Theorem 2.1 from Lemma 3.1.

4 Solving (N_λ) in strongly polynomial time

In this section we devise a strongly polynomial algorithm for solving the primal parametric problem (N_λ) . As before, we assume that a and λ are integral and that the node capacities c are even, so our goal is to find an integer optimal multiflow.

The algorithm starts with computing a (fractional) optimal dual solution l and constructing the BD-graph H w.r.t. the length function $\ell := \bar{a} + \bar{l}$. Then it finds a good IBD-flow f in H saturating the locked edges (assuming validity of Proposition 3.3). Applying the efficient procedure as in the proof of Lemma 3.1 to decompose f into a collection of paths with integer weights, we will obtain an integer optimal solution to (N_λ) .

To provide the desired complexity, we shall work with H given in the *compact* form (defined in Subsection 3.1). The core of our method consists in finding the load function of some integer optimal multiflow F in G (without explicitly computing F itself). This function will just generate the desired IBD-flow in H . We describe the stages of the algorithm in the subsections below.

4.1 Constructing an optimal dual solution

Problem (D_λ) straightforwardly reduces to a “compact” linear program, as follows. Besides variables $l(v) \in \mathbb{Q}_+$ ($v \in VG$), assign a variable $\varphi_s(v) \in \mathbb{Q}$ to each terminal $s \in T$ and node v (a sort of “distance” of v from s). Consider the following problem (where \bar{l} and \bar{a} are defined according to (2.2)):

(4.1) Minimize $c \cdot l$ subject to the following constraints:

$$\begin{aligned} \varphi_s(u) - \varphi_s(v) &\leq \bar{a}(e) + \bar{l}(e) \\ \varphi_s(v) - \varphi_s(u) &\leq \bar{a}(e) + \bar{l}(e) \quad \text{for each } e = uv \in EG; \\ \varphi_s(t) - \varphi_s(s) &\geq \lambda \quad \text{for all } s, t \in T, s \neq t. \end{aligned}$$

Lemma 4.1 *Programs (D_λ) and (4.1) are equivalent.*

Indeed, if (l, φ) is a feasible solution to (4.1) then, obviously, l is a feasible solution to (D_λ) . Conversely, let l be a feasible solution to (D_λ) . For $v \in VG$ and $s \in T$, define $\varphi_s(v) := \text{dist}_\ell(s, v)$, where $\ell := \bar{l} + \bar{a}$. It is easy to check that (l, φ) is a feasible solution to (4.1).

The size of the constraint matrix in (4.1) (written in binary notation) is polynomial in $|VG|$. Therefore, (D_λ) is solvable in strongly polynomial time by use of Tardos’s version of the ellipsoid method. (This remains valid when a and λ are nonnegative rational numbers.)

4.2 Computing the load function of an optimal multiflow

The following fact is of importance.

Lemma 4.2 *One can find, in strongly polynomial time, a function $g: VG \rightarrow \mathbb{Z}_+$ such that $g(v) = \widehat{F}(v)$ for all $v \in VG$, where F is some integer optimal multiflow in (N_λ) .*

Proof. We explain that in order to construct the desired g , it suffices to compare two optimal objective values: one for the original (integer) costs a and the other for certain perturbed costs a_ε . These values are computed by solving the corresponding dual problems by the method described in the previous subsection.

More precisely, let v_1, \dots, v_n be the nodes of G . Let $U := \max_i c(v_i) + 2$, define $\varepsilon(v_i) := 1/U^{i+1}$, $i = 1, \dots, n$, and define the cost function a_ε on VG to be $a + \varepsilon$. Then for any integer feasible multiflow F , we have

$$0 \leq \Phi(F, a, \lambda) - \Phi(F, a_\varepsilon, \lambda) = \sum_i \widehat{F}(v_i) \varepsilon(v_i) < U^{-1} + U^{-2} + \dots + U^{-n} < 1.$$

This and the fact that $\Phi(F, a, \lambda)$ is an integer (as F, a, λ are integral) imply that if F is optimal for a_ε , then F is optimal for a as well. (An integer optimal multiflow for even capacities c exists by Theorem 2.1.) Moreover, for such an optimal F , the number $r := \sum_i \widehat{F}(v_i) \varepsilon(v_i)$ is computed in strongly polynomial time, since it is equal to $c \cdot l - c \cdot l_\varepsilon$, where l and l_ε are optimal dual solutions for a and a_ε , respectively. Here we use the LP duality equalities $\Phi(F, a, \lambda) = c \cdot l$ and $\Phi(F, a_\varepsilon, \lambda) = c \cdot l_\varepsilon$. Also the size of binary encoding of a_ε is bounded by that of a times a polynomial in n , so the dual problem with a_ε is solved in strongly polynomial time w.r.t. the original data.

Hence, we have $rU^{n+1} = \sum_i \widehat{F}(v_i)U^{n-i}$. The number rU^{n+1} is an integer and, in view of $\widehat{F}(v_i) \leq c(v_i) < U$ for each i , the n coefficients in its base U decomposition (the representation via degrees of U) are just $\widehat{F}(v_1), \dots, \widehat{F}(v_n)$, thus giving g . \square

Recall that together with a node load function each multiflow F also induces its edge counterpart (see (1.3)). Lemma 4.2 can be strengthened as follows.

Lemma 4.3 *One can find, in strongly polynomial time, a function $g: VG \cup EG \rightarrow \mathbb{Z}_+$ such that $g(v) = \widehat{F}(v)$ for all $v \in VG$ and $g(e) = \widehat{F}(e)$ for all $e \in EG$, where F is some integer optimal multiflow F in (N_λ) .*

Proof. Split each edge $e = uv$ of G into two edges $ux_e, x_e v$ in series and assign to each new node x_e the capacity $c(x_e) := \min\{c(u), c(v)\}$ and the cost $a(x_e) := 0$. Clearly this transformation does not affect the problem in essence. The node load function, which can be found in strongly polynomial time by Lemma 4.2, yields the desired node and edge load functions in the original graph G . \square

4.3 Constructing an optimal primal solution

Now we explain how to find, in strongly polynomial time, an integer optimal multiflow solving (N_λ) (for a graph G , even node capacities c , rational node costs a , and an integer parameter λ) by using an optimal dual solution l and a function g as in Lemma 4.3. For this g , there exists an integer ℓ -geodesic multiflow F' in G satisfying $F'(v) = g(v)$ for all

$v \in VG$ and $F'(e) = g(e)$ for all $e \in EG$, where $\ell := \bar{a} + \bar{l}$. Our goal is to construct one of such multiflows explicitly.

To do this, we consider the *compact* BD-graph H related to G_ℓ (see Subsection 3.1) and put f to be the function on EH corresponding to g . More precisely, let E' be the subset of edges of H neither incident to q nor contained in the gadgets Γ_w ($w \in V^\sharp$). By the construction of H , there is a natural bijection γ of E' to the set $(VG_\ell - V^\sharp) \cup EG_\ell$. For each $e \in E'$, we set $f(e) := g(\gamma(e))$. In their turn, the values of f on the edges of a gadget Γ_w are assigned as follows: for the loop e_w at θ_w , set $f(e_w) := g(w)$, and for each $s \in T$, set $f(e_{w,s}) := \sum(g(e) : e \in E_{s,w})$, where $E_{s,w}$ is the set of edges in G_ℓ connecting V^s and w . Finally, for each $s \in T$, we set $f(qs) := h(s)$.

Using the fact that the function g on $VG \cup EG$ is determined by some integer optimal (ℓ -geodesic) multiflow F' , it follows that the obtained function f on EH is integer-valued and has zero divergency at all nodes different from q . So f is an IBD-flow in H . Moreover, f is generated by F' as above; in particular, f is good (i.e. satisfies (3.2)). By Lemma 3.1, we can find, in strongly polynomial time, an integer ℓ -geodesic multiflow F generating f . Then F and F' have the coinciding node and edge load functions, and the optimality of F' implies that F is an integer optimal solution to (N_λ) as well, as required.

5 Proof of Proposition 3.3

To complete the proof of Theorem 2.1 it remains to prove Proposition 3.3, which claims the existence of an IBD-flow saturating the “locked” edges. We eliminate the lower capacity constraints (induced by the locked edges) by reducing the claim to the existence of an IBD-flow with a certain prescribed value.

5.1 Maximum IBD-flows

Let Γ be a bidirected graph with a distinguished source q and edge capacities $c: E\Gamma \rightarrow \mathbb{Z}_+$, as described in Subsection 3.2. The classic max-flow min-cut theorem states that the maximum flow value is equal to the minimum cut capacity. A bidirected version of this theorem involves a somewhat more complicated object, called an *odd barrier*. In this subsection we give its definition and state the crucial properties (in Theorems 5.1, 5.2, 5.3) that will be used in the upcoming proof of Proposition 3.3. These properties are nothing else than translations, to the language of bidirected graphs, of corresponding properties established for integer symmetric flows in skew-symmetric graphs, as we will explain in the Appendix.

Let $X \subseteq VG - \{q\}$. The *flip* at (the nodes of) X modifies Γ as follows: for each node $v \in X$ and each edge e incident to v , we reverse the direction of e at v (while preserving the directions of edges at nodes in $VG - X$). For example, if $e = \overrightarrow{u} \overleftarrow{v}$ and $u, v \in X$ then e becomes $\overleftarrow{u} \overrightarrow{v}$, and if $e = \overrightarrow{u} \overrightarrow{v}$ and $u \notin X \ni v$ then e becomes $\overleftarrow{u} \overleftarrow{v}$. BD-graphs Γ and Γ' are called *equivalent* if one is obtained by a flip from the other. Note that flips do not affect bidirected walks in Γ in essence and do not change the maximum value of an IBD-flow in it. We will essentially use flips to simplify requirements in the definition of odd barriers below.

Next, we employ a special notation to designate certain subsets of edges. For $X, Y \subseteq VG$, let $[X, Y]$ denote the set of edges with one endpoint in X and the other in Y . We

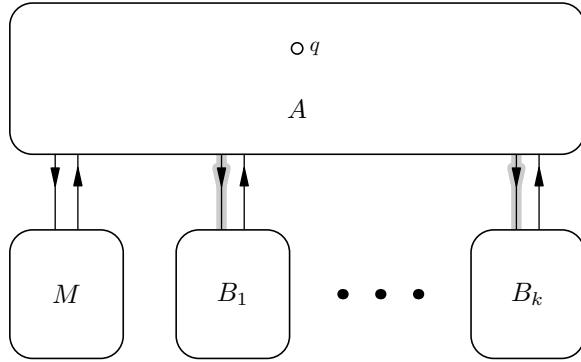


Figure 2: A bidirected odd barrier. Grayed edges correspond to odd capacity constraints (w.r.t. Γ').

will often add arrows above X and/or Y to indicate the subset of edges in $[X, Y]$ directed in one or another way. For example, $[\overleftarrow{X}, \overrightarrow{Y}]$ denotes the set of edges that enter both X and Y , $[\overrightarrow{X}, Y]$ denotes the set of edges leaving X and having the other endpoint in Y (where the direction is arbitrary), and $[\overrightarrow{X}, \overleftarrow{X}]$ denotes the set of edges that leave X at both endpoints (including twice leaving loops). When $Y = V\Gamma - X$, the second term in the brackets may be omitted: $[X]$, $[\overrightarrow{X}]$, and $[\overleftarrow{X}]$ stand for $[X, V\Gamma - X]$, $[\overrightarrow{X}, V\Gamma - X]$, and $[\overleftarrow{X}, V\Gamma - X]$, respectively. Finally, for a function φ on the edges, we write $\varphi[X, Y]$ (rather than $\varphi([X, Y])$) for $\sum_{e \in [X, Y]} \varphi(e)$.

A tuple $\mathcal{B} = (\Gamma' | A, M; B_1, \dots, B_k)$, where Γ' is some BD-graph equivalent to Γ , is called an *odd barrier* for Γ if the following conditions hold with respect to Γ' (see Fig. 2):

- (5.1) (i) A, M, B_1, \dots, B_k give a partition of $V\Gamma' = V\Gamma$, and $q \in A$.
- (ii) For each $i = 1, \dots, k$, $c[\overrightarrow{A}, B_i]$ is odd.
- (iii) For distinct $i, j = 1, \dots, k$, $c[B_i, B_j] = 0$.
- (iv) For each $i = 1, \dots, k$, $c[B_i, M] = 0$.

The *capacity* of \mathcal{B} is defined to be

$$(5.2) \quad c(\mathcal{B}) := 2c[\overrightarrow{A}, \overleftarrow{A}] + c[\overrightarrow{A}] - k.$$

Theorem 5.1 (Max IBD-Flow Min Odd Barrier Theorem) *For Γ, c, q as above, the maximum IBD-flow value is equal to the minimum odd barrier capacity. A (feasible) IBD-flow g and an odd barrier $\mathcal{B} = (\Gamma' | A, M; B_1, \dots, B_k)$ for Γ have maximum value and minimum capacity, respectively, if and only if the following conditions hold with respect to Γ' :*

- (i) $g[\overrightarrow{A}, \overleftarrow{A}] = c[\overrightarrow{A}, \overleftarrow{A}]$ and $g[\overleftarrow{A}, \overrightarrow{A}] = 0$;
- (ii) $g[\overrightarrow{A}, M] = c[\overrightarrow{A}, M]$ and $g[\overleftarrow{A}, M] = 0$;

(iii) for each $i = 1, \dots, k$, either $g[\vec{A}, B_i] = c[\vec{A}, B_i] - 1$ and $g[\overleftarrow{A}, B_i] = 0$, or $g[\vec{A}, B_i] = c[\vec{A}, B_i]$ and $g[\overleftarrow{A}, B_i] = 1$.

Note that there may exist many minimum capacity odd barriers for Γ, c, q . It is well-known that in a usual edge-capacitated digraph with a source s and a sink t , the set of nodes reachable by paths from s in the residual digraph of a maximum s - t flow f determines a minimum capacity s - t cut. Moreover, this minimum cut does not depend on the choice of f and, therefore, may be regarded as the *canonical* one.

A similar phenomenon takes place for maximum IBD-flows and minimum odd barriers (and we will essentially use this in the proof of Proposition 3.3). To describe this, consider an IBD-flow g in Γ . The *residual* BD-graph Γ_g endowed with the *residual* capacity $c_g: E\Gamma_g \rightarrow \mathbb{Z}_+$ is constructed in a similar way as for usual flows. More precisely, $V\Gamma_g = V\Gamma$ and the edges of Γ_g are as follows:

(5.3) (i) each edge $e \in E\Gamma$ with $g(e) < c(e)$ whose residual capacity is defined to be $c_g(e) := c(e) - g(e)$, and
(ii) the *reverse* edge e^R to each edge $e \in E\Gamma$ with $g(e) > 0$; the directions of e^R at the endpoints are reverse to those of e and the residual capacity is $c_g(e^R) := g(e)$.

A bidirected walk P in Γ_g is called c_g -*simple* if P passes each edge e at most $c_g(e)$ times. If P is a c_g -simple closed q - q walk leaving its end q twice, we can increase the value of g in Γ by 2, by sending one unit of flow along P . So the existence (in Γ_g) of such a walk P , which is called (Γ, g) -*residual*, implies that g is not maximum. A converse property holds as well.

Theorem 5.2 *An IBD-flow g in Γ is maximum if and only if there is no (Γ, g) -residual walk.*

When we are given a maximum IBD-flow g , a certain minimum odd barrier can be constructed by considering the residual graph Γ_g . Namely, let $\vec{R}_{\Gamma, g}$ (resp. $\overleftarrow{R}_{\Gamma, g}$) be the set of nodes v that are reachable by a (Γ, g) -residual q - v walk that leaves q and enters v (resp. leaves both q and v). Then $q \notin \vec{R}_{\Gamma, g}$, by the maximality of g .

Theorem 5.3 *Let g be a maximum IBD-flow for Γ, c, q . Define $A := (\vec{R} - \overleftarrow{R}) \cup (\overleftarrow{R} - \vec{R})$ and $M := V\Gamma - (\vec{R} \cup \overleftarrow{R})$, where $\vec{R} := \vec{R}_{\Gamma, g}$ and $\overleftarrow{R} := \overleftarrow{R}_{\Gamma, g}$. Let B_1, \dots, B_k be the node sets of weakly connected components of the underlying undirected subgraph of Γ_g induced by $\vec{R} \cap \overleftarrow{R}$. Define Γ' to be the BD-graph obtained from Γ by flipping the set $\overleftarrow{R} - \vec{R}$ (contained in A). Then $\mathcal{B}_g := (\Gamma' | A, M; B_1, \dots, B_k)$ is a minimum odd barrier.*

An important fact is that the minimum odd barrier \mathcal{B}_f does not depend on g , and we refer to it as the *canonical* odd barrier for Γ, c, q .

Theorem 5.4 *The sets $\overleftarrow{R}_{\Gamma, g}$ are same for all maximum IBD-flows g in Γ , and similarly for the sets $\vec{R}_{\Gamma, g}$, the minimum odd barriers \mathcal{B}_f , and the graphs Γ' obtained from Γ by flipping $\overleftarrow{R}_{\Gamma, g} - \vec{R}_{\Gamma, g}$.*

5.2 Proof of Proposition 3.3

In fact, we have freedom of choosing any of the two forms (expensive or compact) of H to prove Proposition 3.3 in full, as it is easy to see that the claims in these cases are reduced to each other. We prefer to deal with the expensive form, taking advantage from certain nice structural features arising in this case. One reason for our choice is that any IBD-flow in the expensive H is automatically good, as explained in Remark 3.2.

We know that there exists a good fractional bidirected q -flow f in H that saturates the set E_0 of locked edges, and our goal is to show the existence of an IBD-flow saturating E_0 . Recall that any edge $e \in E_0$ is generated by some node v of G , i.e. $e = e_v$.

It will be convenient for us to construct the desired IBD-flow without explicitly imposing the ‘‘lower capacities’’ on the locked edges. For this purpose, we modify H as follows.

First, we add a loop $\overleftarrow{q} \overrightarrow{q}$ with infinite capacity (entering q twice). Also we add to H a node z , which is regarded as a new source.

Second, let E_0 contain an edge $e_v = \overrightarrow{v}^1 \overrightarrow{v}^2$ generated by a vertex $v \in VG_\ell$ in some zone V^s , $s \in T$. We delete e_v from H and, instead, add two edges $\overrightarrow{v}^1 \overleftarrow{z}$ and $\overrightarrow{z} \overrightarrow{v}^2$ with capacity $c(v)$ each.

Third, let E_0 contain the loops e_{w^i} ($i = 1, \dots, c(w)$) for some central node w of G_ℓ . We replace each e_{w^i} (having unit capacity) by edge $\overrightarrow{z} \overrightarrow{\theta}_{w^i}$ with capacity 2; we call it the *root edge* at θ_{w^i} .

We denote the resulting BD-graph by H^1 and keep the previous notation c for its edge capacities. The above q -flow f is transformed, in an obvious way, into a feasible z -flow in H^1 , denoted by f as before. Note that this f saturates all edges created from those in E_0 (i.e. from e_v and e_{w^i} as above); these edges leave z and the value of f is maximum among the feasible z -flows in H^1 and is equal to $2c(E_0)$.

Let g be a maximum IBD-flow in H^1 . We are going to prove that $\text{val}(f) = \text{val}(g)$. This would imply that the corresponding IBD-flow in H saturates E_0 as required. To this aim, consider the *canonical* odd barrier $\mathcal{B} = (H^2 | A, M; B_1, \dots, B_k)$ for H^1, c, z (see Theorem 5.4). Here H^2 is the BD-graph (with the source z) obtained from H^1 according to Theorem 5.3 (i.e. $H^2 := \Gamma'$ for $\Gamma := H^1$). From now on, speaking of edge directions, the capacities c and the flow g , we mean those in H^2 , unless explicitly stated otherwise.

We have (cf. (5.2))

$$(5.4) \quad \text{val}(g) = c(\mathcal{B}) = 2c[\overrightarrow{A}, \overleftarrow{A}] + c[\overrightarrow{A}] - k.$$

The following assertion is crucial.

Lemma 5.5 *For each $p = 1, \dots, k$:*

- (i) $B_p = \{\theta_{w^i}\}$ for some $w \in V^\sharp$ and $i \in \{1, \dots, c(w)\}$;
- (ii) e_{w^i} is not locked (so H^1 contains the loop e_{w^i} but not the root edge at θ_{w^i});
- (iii) among the edges (legs) connecting A and B_p , one edge leaves A and the other edges enter A .

Proof. By the constructions of H and H^1 , for any $w \in V^\natural$ and distinct $i, j = 1, \dots, c(w)$, there is an automorphism $\pi = \pi_{w,i,j}$ of H^1 that swaps θ_{w^i} and θ_{w^j} and is invariant on the other nodes. Also π respects the capacities in H^1 , and the function \tilde{g} induced by g under π (i.e. $\tilde{g}(e) := g(\pi(e))$) is again a maximum IBD-flow in H^1 . Since \mathcal{B} is canonical, it follows from Theorem 5.4 that

$$(5.5) \quad \overrightarrow{R}_{H^1,g} = \overrightarrow{R}_{H^1,\tilde{g}}, \quad \overleftarrow{R}_{H^1,g} = \overleftarrow{R}_{H^1,\tilde{g}}, \quad \overrightarrow{R}_{H^1,g} \cup \overleftarrow{R}_{H^1,g} = A \cup B_1 \cup \dots \cup B_k.$$

The nodes in $\overleftarrow{R}_{H^1,g} - \overrightarrow{R}_{H^1,g}$ are *flipped* when constructing H^2 from H^1 . Then (5.5) and Theorem 5.3 imply that for i, j as above,

- (5.6) (a) θ_{w^i} is flipped if and only if θ_{w^j} is flipped;
- (b) $\theta_{w^i} \in A$ if and only if $\theta_{w^j} \in A$.
- (c) $\theta_{w^i} \in B_1 \cup \dots \cup B_k$ if and only if $\theta_{w^j} \in B_1 \cup \dots \cup B_k$.

Let $p \in \{1, \dots, k\}$. Since the capacity $c[\overrightarrow{A}, B_p]$ (in H^2) is odd, the set $[\overrightarrow{A}, B_p]$ contains an edge e with $c(e)$ odd. Any edge in H^2 having an odd capacity is either a loop e_{w^i} or a leg $e_{w^i,s}$ (regarding “infinite” capacities as even ones).

Obviously, no loop can be “responsible” for the oddness of $c[\overrightarrow{A}, B_p]$.

So $e = e_{w^i,s} = \theta_{w^i}\theta_{w,s}$ for some $w \in V^\natural$, $1 \leq i \leq c(w)$ and $s \in T$. Let \hat{e} denote the edge of H^1 corresponding to e . Then (see Fig. 1(b)) \hat{e} leaves $\theta_{w,s}$ and enters θ_{w^i} . Due to flips, however, this may not be the case for e in H^2 .

Suppose $\theta_{w^i} \in A$ (and $\theta_{w,s} \in B_p$). Then e leaves θ_{w^i} , whence θ_{w^i} is a flipped node in A . Now (5.6)(a,b) imply that all θ_{w^j} are flipped nodes belonging to A and that $e_{w^j,s} \in [\overrightarrow{A}, B_p]$ for all $j = 1, \dots, c(w)$. But then e cannot be “responsible” for the oddness of $c[\overrightarrow{A}, B_p]$ since $c(w)$ is even.

So we have $\theta_{w,s} \in A$ and $\theta_{w^i} \in B_p$. Then e leaves $\theta_{w,s}$. The edge \hat{e} leaves $\theta_{w,s}$ as well. Hence $\theta_{w,s}$ is not flipped. Since $c(w)$ is even, there must be $j \in \{1, \dots, c(w)\}$ such that the leg $e_{w^j,s} = \theta_{w,s}\theta_{w^j}$ is not in $[\overrightarrow{A}, B_p]$ (for otherwise one may pick another pair w', i'). Then θ_{w^j} is not in B_p . In view of (5.6)(c), θ_{w^j} belongs to a B -set in \mathcal{B} different from B_p . Considering the automorphisms $\pi = \pi_{w,i',j'}$ for all distinct $i', j' = 1, \dots, c(w)$ and using the fact that the canonical barrier \mathcal{B} preserves under π (in view of (5.5)), we can conclude that the nodes $\theta_{w^1}, \dots, \theta_{w^{c(w)}}$ belong to different B -sets in \mathcal{B} . Since these B -sets are pairwise disjoint and each automorphism π swaps two copies of θ_w , and do not move the remaining nodes in H^2 , each of these B -sets can contain only a single node. Thus, $B_p = \{\theta_{w^i}\}$, yielding (i) in the lemma.

Next we show (ii). From the construction of H^2 it follows that

$$(5.7) \quad A = \overrightarrow{R}_{H^2,g} - \overleftarrow{R}_{H^2,\tilde{g}} \quad \text{and} \quad B_1 \cup \dots \cup B_k = \overrightarrow{R}_{H^2,g} \cap \overleftarrow{R}_{H^2,g}.$$

By the first equality, any (H^2, g) -residual walk ending at a node $v \in A$ enters v , and by the second equality, there exist an (H^2, g) -residual walk P to $\theta := \theta_{w^i}$ that enters θ and an (H^2, g) -residual walk Q to θ that leaves θ . Recall that the residual walks leave the source z . Let $a = u \xrightarrow{\theta}$ be the last edge of P , and $b = v \xleftarrow{\theta}$ the last edge of Q . Define E' (E'') to be the set of legs $e = e_{w^i,s}$ with $g(e) = 0$ (resp. $g(e) = 1$). Note that (cf. (5.3))

(5.8) if $e \in E'$ then $e^R \notin EH_g^2$ and e enters θ in H_g^2 , and if $e \in E''$ then $e \notin EH_g^2$ and e^R leaves θ in H_g^2 .

Supposing the existence of the root edge $r = \overrightarrow{z\theta}$ (in H^1 and H^2), we can come to a contradiction as follows. Since there is no loop at θ , both nodes u, v are in A . Note that the edge a is different from r (which leaves θ) and from r^R (which enters z). Then (5.8) implies that $a \in E'$. Furthermore, a is of the form $\overrightarrow{u\theta}$. For if a enters u then the edge of P preceding a leaves u , whence the part of P from z to u forms an (H^2, g) -residual walk leaving u , which is impossible since $u \in A$.

So $a \in [\overrightarrow{A}, B_p]$ and $g(a) = 0 = c(a) - 1$. Then $g[\overleftarrow{A}, B_p] = 0$, by Theorem 5.1(iii). This implies that $E'' \subseteq [\overrightarrow{A}, B]$. But then the last edge $b = v\overleftarrow{\theta}$ of the walk Q as above cannot be reverse to any edge in E'' ; for otherwise b enters v , implying that the part of Q from z to v leaves v . Also b is neither reverse to an edge in E' (cf. (5.8)), nor equal to r . The latter is because $r \in [\overrightarrow{A}, B_p]$, and therefore, $g(a) < c(a)$ implies $g(r) = c(r)$ (cf. Theorem 5.1(iii)), whence $r \notin EH_g^2$. Thus, Q does not exist. This contradiction yields (ii).

It remains to show (iii). By (ii), we have $e_{w^i} \in EH^2$, and $[A, B_p]$ is exactly the set of legs at $\theta := \theta_{w^i}$. Suppose $d := |[\overrightarrow{A}, B_p]| \neq 1$. Then $d \geq 3$, since $c[\overrightarrow{A}, B_p] = d$ is odd. Hence $g[\overrightarrow{A}, B_p] \geq d - 1 \geq 2$ (by Theorem 5.1(iii)). Also the fact that all legs enter θ together with $\text{div}_g(\theta) = 0$ and $g(e_{w^i}) \leq 1$ implies that the only possible case is when $g(e_{w^i}) = 1$, $g[\overrightarrow{A}, B_p] = 2$ and $g[\overleftarrow{A}, B_p] = 0$. Now take an (H^2, g) -residual walk Q to θ that leaves θ , and let b be its last edge. Then b is neither the loop e_{w^i} (which is saturated), nor reverse to a leg $e = v\overrightarrow{\theta}$ with $g(e) > 0$. Indeed, if $b = e^R$ then b enters v (in view of $v \in A$ and $e \in [\overrightarrow{A}, B_p]$), and hence the part of Q from z to v leaves v , which is impossible since $v \in A$. This contradiction yields (iii) and completes the proof of the lemma. \square

Based on Lemma 5.5, we now finish the proof of Proposition 3.3. Consider $B_p = \{\theta_{w^i}\}$ and let $e_{w^i, s}$ be the unique edge in $[\overrightarrow{A}, B_p]$. Then $[\overleftarrow{A}, B_p] = \{e_{w^i, t} \mid t \in T - \{s\}\}$. Consider the maximum fractional BD-flow f as before. By the goodness of f (see (3.2)), we have

$$\begin{aligned} f[\overrightarrow{A}, B_p] - f[\overleftarrow{A}, B_p] &= f(e_{w^i, s}) - \sum_{t \in T - \{s\}} f(e_{w^i, t}) \\ &= f(e_{w^i, s}) - (2f(e_{w^i}) - f(e_{w^i, s})) = 2(f(e_{w^i, s}) - f(e_{w^i})) \leq 0 = c[\overrightarrow{A}, B_p] - 1. \end{aligned}$$

Using this and (5.4), we have

$$\begin{aligned} \text{val}(f) &= \text{div}_f(z) = \sum_{v \in A} \text{div}_f(v) = (2f[\overrightarrow{A}, \overleftarrow{A}] - 2f[\overleftarrow{A}, \overrightarrow{A}]) + (f[\overrightarrow{A}] - f[\overleftarrow{A}]) \\ &= (2f[\overrightarrow{A}, \overleftarrow{A}] - 2f[\overleftarrow{A}, \overrightarrow{A}]) + (f[\overrightarrow{A}, M] - f[\overleftarrow{A}, M]) + \sum_{p=1}^k (f[\overrightarrow{A}, B_p] - f[\overleftarrow{A}, B_p]) \\ &\leq 2c[\overrightarrow{A}, \overleftarrow{A}] + c[\overrightarrow{A}, M] + \sum_{p=1}^k (c[\overrightarrow{A}, B_p] - 1) = c[\overrightarrow{A}] + 2c[\overrightarrow{A}, \overleftarrow{A}] - k = c(\mathcal{B}) = \text{val}(g). \end{aligned}$$

Thus, we obtain the desired relation $\text{val}(f) \leq \text{val}(g)$ (which, in fact, holds with equality). This completes the proof of Proposition 3.3.

6 Dual half-integrality

6.1 Polyhedral approach

Theorem 6.1 *Let $a : VG \rightarrow \mathbb{Z}_+$ and $p \in \mathbb{Z}_+$. Then problem (D_λ) has a half-integer optimal solution.*

Proof. The proof follows easily from Theorem 2.1 and the general fact that the “totally dual $1/k$ -integrality” implies the “totally primal $1/k$ -integrality”, which is a natural generalization of a well-known result on TDI-systems due to Edmonds and Giles [EG77]. More precisely, we utilize the following simple fact (see, e.g., [Ka89, Statement 1.1]):

Lemma 6.2 *Let A be a nonnegative $m \times n$ -matrix, b an integral m -vector, and k a positive integer. Suppose that the program $D(c) := \max\{yb \mid y \in \mathbb{Q}_+^m, yA \leq c\}$ has a $1/k$ -integer optimal solution for every nonnegative integral n -vector c such that $D(c)$ has an optimal solution. Then for every nonnegative integral n -vector c , the program $P(c) := \min\{cx \mid x \in \mathbb{Q}_+^n, Ax \geq b\}$ has a $1/k$ -integer optimal solution whenever it has an optimal solution.*

In our case, we set $k := 2$ and take as A (resp. b) the constraint matrix (resp. the right hand side vector) of (D_λ) . Then b is integral, $D(c)$ becomes (N_λ) , $P(c)$ becomes (D_λ) , and the half-integrality for the former implies that for the latter. \square

This proof is not “constructive” and does not lead directly to an efficient method for finding a half-integer optimal solution l to (D_λ) . Below we devise a strongly polynomial algorithm.

6.2 The algorithm

It has as the input arbitrary (rational-valued) optimal solutions l and F to (D_λ) and (N_λ) , respectively, and outputs a half-integer optimal solution \hat{l} to (D_λ) . (Such l and F can be found in strongly polynomial time as described in Section 4.)

As before, we set $\ell := \bar{a} + \bar{l}$, and in what follows, speaking of a geodesic, we always mean an ℓ -geodesic in G , i.e. a T -path P with $\ell(P) = a(P) + l(P) = \lambda$. Our goal is to construct $\hat{l} : VG \rightarrow \frac{1}{2}\mathbb{Z}_+$ satisfying the following conditions:

- (6.1) (i) $a(P) + \hat{l}(P) \geq \lambda$ for any T -path P in G ;
- (ii) $a(P) + \hat{l}(P) = \lambda$ for each geodesic P ;
- (iii) for $v \in VG$, if $l(v) = 0$ then $\hat{l}(v) = 0$.

Then (6.1) and the complementary slackness conditions (2.6)–(2.7) imply that the node length \hat{l} forms an optimal solution to (D_λ) .

We construct an undirected graph Γ and endow it with edge lengths $\mu : E\Gamma \rightarrow \mathbb{Z}_+$ as follows. We first include in Γ the terminal set T and all nodes and edges of G contained in geodesics. Also we add to Γ the edges of G with both ends lying on geodesics or T . The edges e of the current Γ are called *regular* and we define $\mu(e) := 0$.

Next we add to Γ additional edges, which are related to constraints due to parts of G outside Γ . More precisely, we scan all pairs of nodes $u, v \in VT$ not connected by a

(regular) edge and such that there exists a path Q in G having all nodes in $VG - V\Gamma$ and whose first node is adjacent to u , and the last node to v . We add to Γ edge $e = uv$, referring to it as a *virtual* edge, and define its length $\mu(e)$ to be the minimum value of $a(Q)$ among such paths Q . The construction of Γ, μ reduces to a polynomial number of usual shortest paths problems in G .

Note that $l(v) = 0$ holds for each node $v \in VG - V\Gamma$ (by (2.6) and (2.7)). We assign $\widehat{l}(v) := 0$ for these nodes v and will further focus on finding values of \widehat{l} on the nodes in Γ .

For a path P in Γ , let $\Lambda(P)$ denote its *full length* $a(P) + l(P) + \mu(P)$. Clearly $\Lambda(P) \geq \lambda$ holds for any T -path P in Γ , and for each T -path Q in G , there exists a *shortcut* path P in Γ such that $\Lambda(P) \leq a(Q) + l(Q)$.

The desired lengths \widehat{l} on $V\Gamma$ will be extracted from a system of linear constraints described below. For a node $v \in V\Gamma$, let T_v (Π_v) denote the set of terminals $s \in T$ (resp. pairs $s, t \in T$) such that v belongs to a geodesic from s (resp. connecting s and t). When a terminal s belongs to no geodesic, we set by definition $T_s := \{s\}$. For each $v \in V\Gamma$ and $s \in T_v$, we introduce two variables $\rho_s^-(v)$ and $\rho_s^+(v)$ and impose the following constraints:

(6.2) (i) For each $s \in T$, $\rho_s^-(s) = 0$.

(ii) For each $v \in V\Gamma$ and $s \in T_v$,

$$\begin{aligned} \rho_s^+(v) - \rho_s^-(v) &= a(v) && \text{if } l(v) = 0, \\ &\geq a(v) && \text{if } l(v) > 0. \end{aligned}$$

(iii) For each $v \in V\Gamma$ and $\{s, t\} \in \Pi_v$, $\rho_s^+(v) + \rho_t^-(v) = \lambda$ (and $\rho_t^+(v) + \rho_s^-(v) = \lambda$).

(iv) If $e = uv \in E\Gamma$ and $s \in T_u \cap T_v$, then

$$\begin{aligned} \rho_s^-(v) - \rho_s^+(u) &\leq \mu(e), \\ \rho_s^-(u) - \rho_s^+(v) &\leq \mu(e). \end{aligned}$$

Moreover, if there exists a geodesic from s containing both u, v in this order (resp. in the order v, u), then the former (resp. the latter) inequality is replaced by equality. (Note that in this case $\mu(e) = 0$.)

(v) If $e = uv \in E\Gamma$, $s \in T_u$, $t \in T_v$, and $s \neq t$, then $\rho_s^+(u) + \rho_t^+(v) \geq \lambda - \mu(e)$.

The meaning of these variables becomes evident from the proof of the next statement.

Lemma 6.3 *System (6.2) has a solution.*

Proof. For a $p-q$ path P in Γ , define its *pre-length* to be $\Lambda(P) - (a(q) + l(q))$ (i.e. compared with the full length, we do not count the last node). For $v \in V\Gamma$ and $s \in T_v$, define $\rho_s^-(v)$ (resp. $\rho_s^+(v)$) to be the minimum pre-length (resp. the minimum full length) of an $s-v$ path in Γ . Then (6.2)(i)–(iii) follow from the construction. Condition (6.2)(iv) represents a sort of triangle inequalities (giving one equality if e belongs to a geodesic from s). Finally, condition (6.2)(v) holds since the full length of any $s-t$ path in Γ is at least λ . \square

We observe that in linear system (6.2), each constraint contains at most two variables, each occurring with the coefficient 1 or -1 , and that the R.H.S. in it is an integer. A well-known fact is that a linear system with such features is totally dual half-integral; therefore, it has a half-integer basis solution (whenever it has a solution at all), and such a solution can be found in strongly polynomial time (cf., e.g., [EJ70, Sc03]).

Given a half-integer solution (ρ^-, ρ^+) to (6.2), we define half-integer node lengths \hat{l} as follows:

$$\hat{l}(v) := \rho_s^+(v) - \rho_s^-(v) - a(v) \quad \text{for all } v \in V\Gamma \text{ and } s \in T_v.$$

Now the desired algorithmic result is provided by the following

Lemma 6.4 \hat{l} is well-defined and satisfies (6.1).

Proof. We first show that for any $v \in V\Gamma$ and $s, t \in T_v$,

$$(6.3) \quad \rho_s^+(v) - \rho_s^-(v) = \rho_t^+(v) - \rho_t^-(v).$$

This is trivial when $\Pi_v = \emptyset$ (since in this case $v \in T$ and $T_v = \{v\}$). Let $\Pi_v \neq \emptyset$. If $\{s, t\} \in \Pi_v$, then (6.3) follows from (6.2)(iii). Now (6.3) with any two $s, t \in T_v$ is implied by the fact that the graph whose nodes and edges are the elements of T_v and Π_v , respectively, is connected (as it is easy to see that for $\{s, t\}, \{p, q\} \in \Pi_v$, at least one of $\{s, p\}, \{s, q\}$ is in Π_v as well). So \hat{l} is well-defined.

Property (6.1)(iii) is immediate from (6.2)(ii).

To see (6.1)(ii), consider an s - t geodesic P . Going along P step by step and applying (6.2)(ii),(iv), we observe that for each node v on P , the s - v part P' of P satisfies $a(P') + \hat{l}(P') = \rho_s^+(v) - \rho_s^-(s)$. When reaching t , we obtain $a(P) + \hat{l}(P) = \rho_s^+(t) - \rho_s^-(s)$, and now (6.1)(ii) follows from (6.2)(iii) and $\rho_s^-(s) = \rho_t^-(t) = 0$ (by (6.2)(i)).

Finally, consider an arbitrary T -path Q in Γ , from p to q say. To conclude with (6.1)(i), it suffices to show that

$$(6.4) \quad \Delta(Q) := a(Q) + \hat{l}(Q) + \mu(Q) \geq \lambda.$$

Represent Q as the concatenation $Q' \cdot Q''$, where Q' is a part of a geodesic from p . We prove (6.4) by induction on the number $|Q''|$ of edges in Q'' . When $|Q''| = 0$, Q is a geodesic, and we are done. Assuming this is not the case, take the first edge $e = uv$ of Q'' , where u is the end of Q' . By reasonings above, $\Delta(Q') = \rho_p^+(u)$. If $v \in T$ (and therefore, $v = q$), (6.4) immediately follows from (6.2)(v) (with $s := q$ and $t := q$). And if $v \notin T$, then v belongs to some s - t geodesic L . W.l.o.g., one may assume that $s \neq p$ and $t \neq q$. Applying (6.2)(v) to s, p, e , we have

$$\rho_p^+(u) + \rho_s^+(v) + \mu(e) \geq \lambda.$$

Comparing this with $\rho_s^-(v) + \rho_t^+(v) = \lambda$ and using $\rho_s^+(v) - \rho_s^-(v) = a(v) + \hat{l}(v)$, one can conclude that $\Delta(Q) \geq \Delta(R)$, where R is the t - q path being the concatenation of the t - v part of (the reverse of) L and the v - q part R'' of Q . Since $|R''| = |Q''| - 1$, we can apply induction and obtain $\Delta(Q) \geq \Delta(R) \geq \lambda$, as required. \square

Appendix: Skew-symmetric graphs and flows

In this section we recall the notions of skew-symmetric graphs and integer skew-symmetric flows, review known results on such graphs and flows, and then use them to derive necessary results on bidirected graphs and flows to which we appealed in Section 5.

7.1 Skew-symmetric graphs

A *skew-symmetric graph*, or an *SK-graph* for short, is a digraph $G = (VG, EG)$, with possible parallel arcs, endowed with two bijections σ_V, σ_A such that: σ_V is an involution on the nodes (i.e. $\sigma_V(v) \neq v$ and $\sigma_V(\sigma_V(v)) = v$ for each node v); σ_A is an involution on the arcs; and for each arc a from u to v , $\sigma_A(a)$ is an arc from $\sigma_V(v)$ to $\sigma_V(u)$. For relevant results on SK-graphs and a relationship between SK- and BD-graphs, see [Tu67, GK96, GK04, BK07]. For brevity σ_V and σ_A are combined into one mapping σ on $VG \cup AG$, which is called the *symmetry* (or skew-symmetry, to be precise) of G . For a node (arc) x , its symmetric node (arc) $\sigma(x)$ is also called the *mate* of x , and we usually use notation with primes for mates, denoting $\sigma(x)$ by x' . Although G is allowed to contain parallel arcs, when it is not confusing, an arc from u to v may be denoted as (u, v) or \overrightarrow{uv} .

Observe that if G contains an arc a from a node v to its mate v' , then a' is also an arc from v to v' (i.e. a' is parallel to a).

The symmetry σ is extended in a natural way to walks, subgraphs and other objects in G . In particular, two walks are symmetric to each other if the elements of one of them are symmetric to those of the other and go in the reverse order: for a walk $P = (v_0, a_1, v_1, \dots, a_k, v_k)$, the symmetric walk $P' = \sigma(P)$ is $(v'_k, a'_k, v'_{k-1}, \dots, a'_1, v'_0)$.

Next we explain a relationship between skew-symmetric and bidirected graphs. Given an SK-graph G , choose an arbitrary partition $\pi = (V_1, V_2)$ of VG such that $V_2 = \sigma(V_1)$. Then G and π determine the BD-graph G^* with $VG^* = V_1$ whose edges correspond to the pairs of symmetric arcs in G . More precisely, arc mates a, a' of G generate one edge e of G^* connecting nodes $u, v \in V_1$ such that: (i) e goes from u to v if one of a, a' goes from u to v (and the other goes from v' to u' in V_2); (ii) e leaves both u, v if one of a, a' goes from u to v' (and the other from v to u'); (iii) e enters both u, v if one of a, a' goes from u' to v (and the other from v' to u). Note that e becomes a loop if a, a' connect a pair of symmetric nodes.

Conversely, a BD-graph G^* determines an SK-graph G with symmetry σ as follows. Make a copy $\sigma(v)$ of each element v of $V^* := VG^*$, forming the set $(V^*)' := \{\sigma(v) \mid v \in V^*\}$. Put $VG := V^* \sqcup (V^*)'$. For each edge e of G^* connecting nodes u and v , assign two “symmetric” arcs a, a' in G so as to satisfy (i)–(iii) above (where $u' = \sigma(u)$ and $v' = \sigma(v)$). An example is depicted in Fig. 3.

Remark 7.1 Note that one BD-graph generates one SK-graph, by the second construction. On the other hand, one SK-graph generates a set of BD-graphs, depending on the partition π of V , by the first construction. Namely, for each pair of symmetric mates $\{v, v'\}$ in G one may distribute v, v' between V_1, V_2 so that either $v \in V_1, v' \in V_2$ or, reversely, $v \in V_2, v' \in V_1$. The resulting BD-graphs are obtained from one other by making corresponding flips (defined in Subsection 5.1).

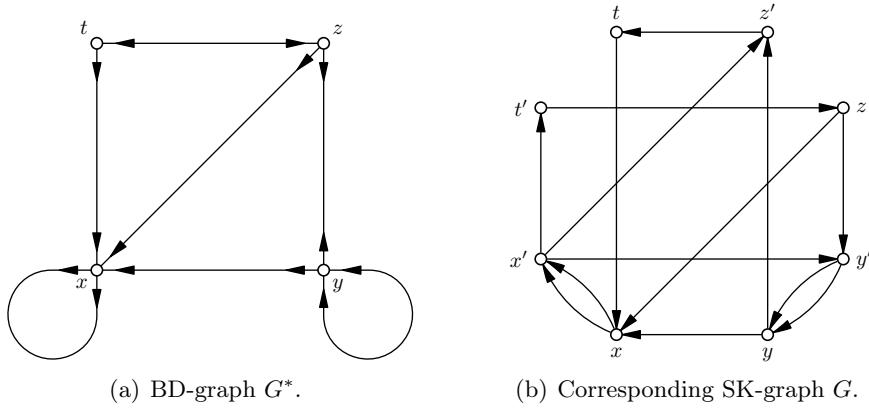


Figure 3: Related bidirected and skew-symmetric graphs.

There is essentially a one-to-one correspondence between the walks in G^* and G . More precisely, let τ be the natural mapping of $VG \cup AG$ to $VG^* \cup EG^*$. Each walk $P = (v_0, a_1, v_1, \dots, a_k, v_k)$ in G (where $a_i = (v_{i-1}, v_i)$) induces the sequence

$$\tau(P) := (\tau(v_0), \tau(a_1), \tau(v_1), \dots, \tau(a_k), \tau(v_k))$$

of nodes and edges in G^* . One can see that $\tau(P)$ is a walk in G^* (i.e. $\tau(a_i), \tau(a_{i+1})$ form a transit pair at $\tau(v_i)$, for each i) and that $\tau(P')$ is the walk reverse to $\tau(P)$. Moreover, for any walk P^* in G^* , there is exactly one walk P in G such that $\tau(P) = P^*$ (considering P up to replacing an arc $a \in AP$ by its mate a' when a, a' are parallel, i.e. correspond to a loop in G^*).

7.2 Skew-symmetric flows

We call a function φ on the arcs of an SK-graph G (*self*-)*symmetric* if $\varphi(a) = \varphi(a')$ for all $a \in AG$. Let $s \in VG$ be a designated *source*; its mate s' is regarded as the *sink*. An *integer skew-symmetric $s-s'$ flow*, or an *ISK-flow* for short, is a symmetric function $f: AG \rightarrow \mathbb{Z}_+$ being an $s-s'$ flow in a usual sense: $\text{div}_f(v) = 0$ for all $v \in VG - \{s, s'\}$, and $\text{div}_f(s) \geq 0$. The *value* of f is $\text{val}(f) := \text{div}_f(s)$. Here $\text{div}_f(v)$ denotes the usual divergence (given by (3.1), where $\delta^{\text{in}}(v)$ and $\delta^{\text{out}}(v)$ are the sets of arcs entering and leaving v , respectively).

For a capacity function $c: AG \rightarrow \mathbb{Z}_+$, a flow f is said to be *feasible* if $f(a) \leq c(a)$ for all $a \in AG$. We refer to a feasible ISK-flow of maximum possible value as a *maximum ISK-flow*.

The above correspondence between BD- and SK-graphs is naturally extended to flows. More precisely, if f is a symmetric $s-s'$ flow in G , then transferring the values of f from the pairs of arc mates of G to the edges of the BD-graph $G^* := \tau(G)$, we obtain a $\tau(s)$ -flow in G^* , denoted as f^* . The converse correspondence is evident as well.

For $X, Y \subseteq VG$, let (X, Y) denote the set of arcs going from X to Y . Also (accommodating notation from Section 5 to digraphs) we denote by $[\vec{X}]$ the set of arcs leaving X .

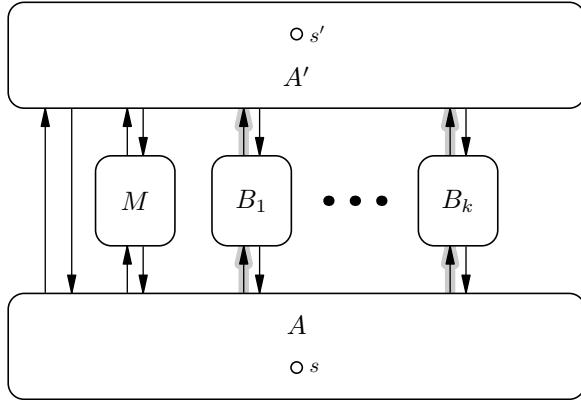


Figure 4: A skew-symmetric odd-barrier. Grayed arcs correspond to odd capacity constraints.

Let $c : AG \rightarrow \mathbb{Z}_+$ be a symmetric capacity function. Then a tuple $\mathcal{B} = (A, M; B_1, \dots, B_k)$ of subsets of VG is called a (skew-symmetric) *odd barrier* (w.r.t. the source s) if the following conditions hold (see Fig. 4):

(7.1) (i) the sets $A, A' = \sigma(A), M, B_1, \dots, B_k$ give a partition of VG , each B_i is self-symmetric ($B'_i = B_i$), and $s \in A$;
(ii) For each i , $c(A, B_i)$ is odd.
(iii) For distinct i, j , $c(B_i, B_j) = 0$.
(iv) For each i , $c(B_i, M) = c(M, B_i) = 0$.

The *capacity* of \mathcal{B} is defined to be

$$(7.2) \quad c(\mathcal{B}) := c[\vec{A}] - k.$$

Odd barriers in skew-symmetric graphs are related to their bidirected counterparts introduced in Section 5. Indeed, consider a BD-graph G^* with integer edge capacities $c : EG^* \rightarrow \mathbb{Z}_+$ and a source s . Construct the related SK-graph G with $VG = V \sqcup V'$, where $V := VG^*$. Edge capacities c in G^* induce symmetric arc capacities in G , also denoted by c . The source s in G^* gives the source s and the sink s' in G . Consider a skew-symmetric odd barrier $\mathcal{B} = (A, M; B_1, \dots, B_k)$ in G .

This barrier gives rise to the following odd BD-barrier \mathcal{B}^* in G^* obeying $c(\mathcal{B}^*) = c(\mathcal{B})$. We first construct a new BD-graph from G by taking a bipartition $(V_1, V_2 = \sigma(V_1))$ of VG such that $A \subseteq V_1$ and $V_1 - (A \cup A') = V - (A \cup A')$; cf. Remark 7.1. The resulting BD-graph H^* is equivalent to G^* . Moreover, H^* is obtained from G^* by flipping a subset of nodes within A .

The node subsets M, B_1, \dots, B_k in G are self-symmetric and induce subsets M^*, B_1^*, \dots, B_k^* in G^* and H^* in a natural way; namely, $M^* := M \cap V = M \cap V_1$ and similarly for B_i^* . Define $\mathcal{B}^* := (H^*|A^*, M^*; B_1^*, \dots, B_k^*)$, where $A^* := (A \cup A') \cap V$. A straightforward examination shows that the properties in (7.1) imply their bidirected counterparts in (5.1). To see that $c(\mathcal{B}^*) = c(\mathcal{B})$, define $Z := M \cup B_1 \cup \dots \cup B_k$. Note that $c[\vec{A}] = c(A, A') + c(A, Z)$. The capacity $c(A, A')$ is equal to $2c[\vec{A}^*, \vec{A}^*]$ (in H^*) since

(A, A') consists of pairs of arc mates, each pair corresponding to an edge in $[\overrightarrow{A^*}, \overleftarrow{A^*}]$. And the capacity $c(A, Z)$ is equal to $c[\overrightarrow{A^*}]$ since the (symmetric) set Z corresponds to $M^* \cup B_1^* \dots B_k^*$ in H^* .

In light of these observations, Theorem 5.1 is a consequence of the following Tutte's theorem. (For shorter proofs of this and next theorems, see also [GK04].)

Theorem 7.2 (Max ISK-Flow Min Odd Barrier Theorem [Tu67]) *For G, c, s as above, the maximum ISK-flow value is equal to the minimum odd barrier capacity. An ISK-flow f and an odd barrier $\mathcal{B} = (A, M; B_1, \dots, B_k)$ have maximum value and minimum capacity, respectively, if and only if the following hold:*

- (i) $f(A, A' \cup M) = c(A, A' \cup M)$ and $f(A' \cup M, A) = 0$;
- (ii) for each $i = 1, \dots, k$, either $f(A, B_i) = c(A, B_i) - 1$ and $f(B_i, A) = 0$, or $f(A, B_i) = c(A, B_i)$ and $f(B_i, A) = 1$.

Next we establish additional correspondences. Consider an ISK-flow f in G . The *residual* SK-graph G_f endowed with the *residual* capacities $c_f: AG \rightarrow \mathbb{Z}_+$ is constructed in a standard fashion: $VG_f = VG$, and the arcs of G_f are:

(7.3) (i) each arc $a \in AG$ with $f(a) < c(a)$ whose residual capacity is defined to be $c_f(a) := c(a) - f(a)$, and
(ii) the *reverse* arc $a^R = (v, u)$ to each arc $a = (u, v) \in AG$ with $f(a) > 0$; its residual capacity is $c_f(a^R) := f(a)$

(cf. (5.3)). A path P in G_f is called c_f -*regular* if $c_f(a) = c_f(a') \geq 2$ holds for each pair of arc mates a, a' occurring in P . (In other words, the bidirected image of P in G_f^* is a c_f -simple walk.) If P is a c_f -regular $s-s'$ path, we can increase the value of f by 2 (by sending one unit of flow along P and one unit of flow along P'). So the existence of such a P implies the non-maximality of f . A converse property is valid as well.

Theorem 7.3 ([Tu67]) *An ISK-flow f is maximum if and only if there is no c_f -regular $s-s'$ path in G_f .*

This implies Theorem 5.2 for IBD-flows.

Given a maximum ISK-flow f , a certain minimum odd barrier can be constructed by considering the residual graph G_f . The construction described in the proof of Theorem 3.5 in [GK04] (relying on Lemma 2.2 in [GK96]) is as follows.

Theorem 7.4 *Let f be a maximum ISK-flow. Let $R = R_f$ be the set of nodes reachable from s by c_f -regular paths in G_f . Define $A := R - R'$ and $M := VG - (R \cup R')$. Let B_1, \dots, B_k be the node sets of weakly connected components of the subgraph G induced by $R \cap R'$. Then $\mathcal{B}_f := (A, M; B_1, \dots, B_k)$ is a minimum odd barrier. \square*

This subset R of nodes in G corresponds to two sets $\overrightarrow{R} = \overrightarrow{R}_{G^*, f^*}$ and $\overleftarrow{R} = \overleftarrow{R}_{G^*, f^*}$ in G^* (defined just before Theorem 5.3; here $\Gamma = G^*$ and $g = f^*$). More precisely, assuming that each node $v \in VG^*$ corresponds to node mates v, v' in G (cf. Subsection 7.1), one can realize that \overrightarrow{R} (resp. \overleftarrow{R}) is the set of nodes $v \in VG^*$ such that $v \in R$ (resp. $v' \in R$).

Finally, the last theorem in Subsection 5.1 (Theorem 5.4) is implied by the following assertion.

Theorem 7.5 *The sets R_f in Theorem 7.4 are equal for all maximum ISK-flows f . Therefore, the minimum odd barriers \mathcal{B}_f are equal as well.*

This fact can be extracted from reasonings in [GK04], yet it is not formulated there explicitly. For this reason, we give a direct proof.

Let f be a maximum ISK-flow such that the set R_f is inclusion-wise minimal and let $\mathcal{B}_f := (A, M; B_1, \dots, B_k)$. Consider another maximum ISK-flow g (if any). Then

$$(7.4) \quad c[\vec{A}] - k = c(\mathcal{B}) = \text{val}(g) = \sum_{v \in A} \text{div}_g(v) = g\langle A, A' \rangle + g\langle A, M \rangle + g\langle A, B_1 \rangle + \dots + g\langle A, B_k \rangle,$$

where for disjoint subsets $X, Y \subset VG$, $g\langle X, Y \rangle$ denotes $g(X, Y) - g(Y, X)$. For $i = 1, \dots, k$, we have: $g(A, B_i) \leq c(A, B_i)$; $c(A, B_i)$ is odd; and $g\langle A, B_i \rangle$ is even (the latter is due to a result in [Tu67]; see also [GK04, Corollary 3.2]). Therefore, $c(A, B_i) - g\langle A, B_i \rangle \geq 1$. Also $g\langle A, A' \rangle \leq c(A, A')$ and $g\langle A, M \rangle \leq c(A, M)$. Comparing these relations with (7.4), we conclude that:

- (i) all arcs in $(A, A' \cup M)$ are saturated by g , while all arcs a in $(A' \cup M, A)$ are free of g (i.e. $g(a) = 0$);
- (ii) for each i , $g\langle A, B_i \rangle = c(A, B_i) - 1$.

In terms of the residual graph G_g , (i) and (ii) mean that the sets A and $VG - A$ are connected in G_g by exactly k arcs a_1, \dots, a_k , each a_i going from A to B_i and having the residual capacity 1. By symmetry, A' and $VG - A'$ are connected by only the arcs a'_1, \dots, a'_k (each a'_i goes from B_i to A' , and $c_g(a_i) = 1$). Also by (7.1)(iii),(iv), no arc in G_g connects different sets among M, B_1, \dots, B_k . Therefore, the set R_g of nodes in G_g reachable from s by c_g -regular paths is contained in R_f . By the minimality of R_f , we have $R_f = R_g$, as required. \square

Acknowledgements. We thank the anonymous referee for useful suggestions.

References

- [BK07] M.A. Babenko and A.V. Karzanov. Free multiflows in bidirected and skew-symmetric graphs. *Discrete Appl. Math.*, 155(13):1715–1730, 2007.
- [Ch77] B.V. Cherkassky. A solution of a problem on multicommodity flows in a network. *Ekonomika i Matematicheskie Metody*, 13(1):143–151, 1977; in Russian.
- [EG77] J. Edmonds and R. Giles. A min–max relation for submodular functions on graphs. *Ann. Discrete Math.*, 1:185–204, 1977.
- [EJ70] J. Edmonds and E. L. Johnson. Matching: a well-solved class of integer linear programs. In R. Guy, H. Hanani, N. Sauer, and J. Schönhein, eds, *Combinatorial Structures and Their Applications*, Gordon and Breach, NY, 1970, pp. 89–92.
- [FF62] L. Ford and D. Fulkerson. *Flows in Networks*. Princeton University Press, 1962.

[GK96] A.V. Goldberg and A.V. Karzanov. Path problems in skew-symmetric graphs. *Combinatorica*, 16(3):353–382, 1996.

[GK97] A.V. Goldberg and A.V. Karzanov. Scaling methods for finding a maximum free multiflow of minimum cost. *Math. Oper. Res.*, 22(1):90–109, 1997.

[GK04] A.V. Goldberg and A.V. Karzanov. Maximum skew-symmetric flows and matchings. *Math. Program.*, 100(3):537–568, 2004.

[IKN98] T. Ibaraki, A.V. Karzanov, and H. Nagamochi. A fast algorithm for finding a maximum free multiflow in an inner Eulerian network and some generalizations. *Combinatorica*, 18(1):61–83, 1998.

[Ka79] A.V. Karzanov. A problem on maximum multiflow of minimum cost. In: *Combinatorial Methods for Flow Problems* (Inst. for System Studies Press, Moscow, 1979, issue 3), pp. 138–156, in Russian.

[Ka89] A.V. Karzanov. Polyhedra related to undirected multicommodity flows. *Linear Algebra and its Applications*, 114/115:293–328, 1989.

[Ka94] A.V. Karzanov. Minimum cost multiflows in undirected networks. *Math. Program.*, 66(3):313–325, 1994.

[La76] E.L. Lawler. *Combinatorial Optimization: Networks and Matroids*. Holt, Reinhardt, and Winston, NY, 1976.

[Lo76] L. Lovász. On some connectivity properties of eulerian graphs. *Acta Math. Akad. Sci. Hung.*, 28:129–138, 1976.

[Pa07] G. Pap. Some new results on node-capacitated packing of a-paths. In *STOC '07: Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, pages 599–604, New York, NY, USA, 2007. ACM.

[Sc03] A. Schrijver. *Combinatorial Optimization*. Springer, Berlin, 2003.

[Tu67] W.T. Tutte. Antisymmetrical digraphs. *Canadian J. Math.*, 19:1101–1117, 1967.

[Va01] V. Vazirani. *Approximation Algorithms*. Springer, 2001.